Least-Squares Estimation of Group Delay for Astrometric Interferometers

Peter R. Lawson, M. Mark Colavita, and Philip J. Dumont

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109

ABSTRACT

The Palomar Testbed Interferometer is an astrometric interferometer that uses both phase and group-delay measurements for narrow-angle astrometry. The group delay measurements are performed using 5 spectral channels across the band from 2.0 to 2.4 microns. Group delay is estimated from phasors (sine and cosine of fringes) calculated at each spectral channel using pathlength modulation of 1 wavelength. Normally the group-delay is estimated to be the delay corresponding to the peak of the power spectrum of the complex Fourier transform of these phasors. The Fourier transform does not however yield a least-squares estimate of the delay. It is natural to suppose that the precision of phase estimation could be achieved in a group-delay estimate using a least-squares approach. We describe the least-squares group-delay estimator that has been implemented at PTI and illustrate its performance as applied to narrow-angle astrometry.

Keywords: astronomy, astrometry, interferometry, group delay

1. INTRODUCTION

The Palomar Testbed Interferometer (PTI) is an infrared astrometric interferometer located at the Palomar Observatory near San Diego in southern California. It was developed at the Jet Propulsion Laboratory to test interferometric techniques applicable to the Keck Interferometer and other interferometers within NASA's Origins program. PTI was designed as a testbed for methods of narrow-angle astrometry and phase-referencing, techniques that may be used for the direct detection of extra-solar planets in future space missions.

Astrometric interferometers measure delay-line position of two or more stars and added reference stars to estimate the baseline orientation and angular separation of objects in the sky. Phase measurements techniques are invariably used because they potentially offer the highest resolution for delay estimation. A zero-seeking fringe tracking servo is used at PTI to keep the measured phase at zero. The expected and rms phase error $\sigma_{\phi d}$ for a four-bin phase estimate is then Wyant (1975):

$$\sigma_{\phi d} = \frac{\pi}{2V\sqrt{N}}, \tag{1}$$

The rms error in *delay* inferred from a phase measurement $\sigma_{\phi d}$ using a white-light fringe of central wavelength λ_{wl} is therefore

$$\sigma_{\phi d}\,\frac{\lambda_{wl}}{2\pi} = \frac{\lambda_{wl}}{4V\sqrt{N}}, \tag{2}$$

where V is the fringe visibility, N is the number of photons per bin, and the effects of background and detector read-noise have been ignored. The oft-quoted drawback of phase-estimation is the 2π ambiguety

Further author information: Send correspondence to PRL
E-mail: lawson@huey.jpl.nasa.gov

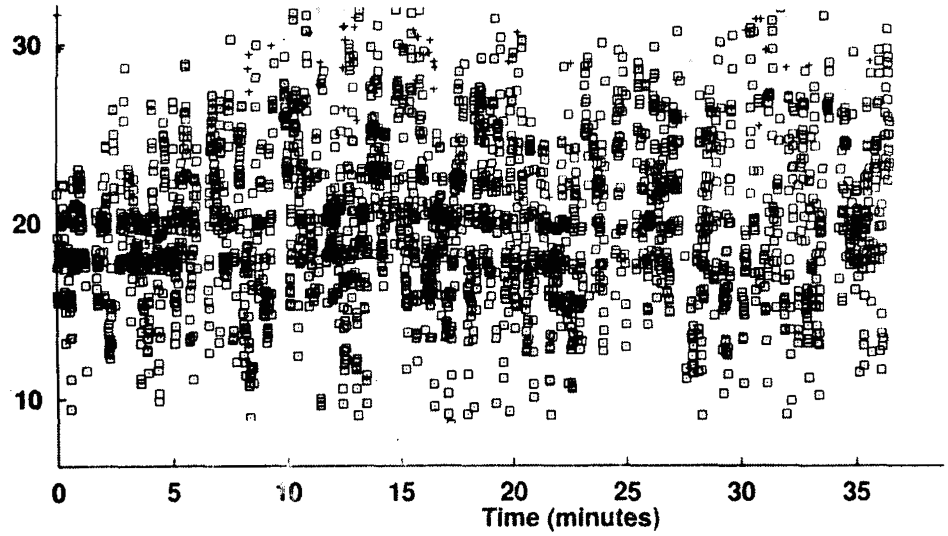

Figure 1. Group-delay as a function of time showing fringe hops. An example under relavitely poor seeing conditions at the Palomar Testbed Interferometer.

in phase; such that even though the rms phase error may be small, the fringe tracker may occasionally (or perhaps frequently) change its tracking location by hopping in delay one wavelength, from the zero-phase location on fringe to another. If the seeing conditions are poor, it may be difficult to estimate the location of the central fringe in the fringe packet, and the fringe tracker will erroneously divide its time between the brightest fringes near the peak of the coherence envelope. This behavior is illustrated in Figure 1.

The two modern astrometric interferometers presently in operation, the Navy Prototype Optical Interferometer (NPOI) and the Palomar Testbed Interferometer (PTI), therefore use group delay estimation to unwrap phase estimates, to unambiguously determine the central fringe in the fringe packet. In each case, although the details are somewhat different, the group delay estimate is derived by measuring fringe phasors at several different wavelengths simultaneously and Fourier transforming the phasors. The power spectrum thus derived will be featureless except for a peak at the spatial frequency of the fringe (corresponding to the delay) and random noise. For example, at PTI a 5-pixel spectrometer is used that samples wavelengths from 2.0 to 2.4 μm. At each wavelength a four-bin phase estimate is used to derive phasors (the sine and cosine of the phase) that are then Fourier transformed.

The variance of a group-delay estimate is typically much worse that of a phase estimate. A delay measurement by power-spectrum analysis can be thought of as an excersise in optimal estimation, or function fitting in the presence of noise. The signal-to-noise ratio of V^2 measurements at a delay s can be written $\mathrm{SNR}(s) = |\Lambda(s)|^2/\sigma_s$, where $\Lambda(s)$ is the amplitude of the Fourier transform of the fringe, and σ_s is the rms noise in the power spectrum. Ideally the response in the power spectrum is a sinc function whose width is inversely proportional to the total bandwidth of the spectrometer $\Delta\kappa$ and whose height Λ is proportional to the SNR of the measurements. The sinc function may be written

$$\frac{\sin \pi x \Delta\kappa}{\pi x \Delta\kappa}, \tag{3}$$

where a standard fast Fourier transform (without any padding) would produce a power spectrum with samples separated in delay by $\Delta x = 1/\Delta\kappa$. The rms uncertainty in delay σ_{gd} is therefore

$$\sigma_{gd} = \sigma_v \frac{1/\Delta\kappa}{\Lambda}, \tag{4}$$

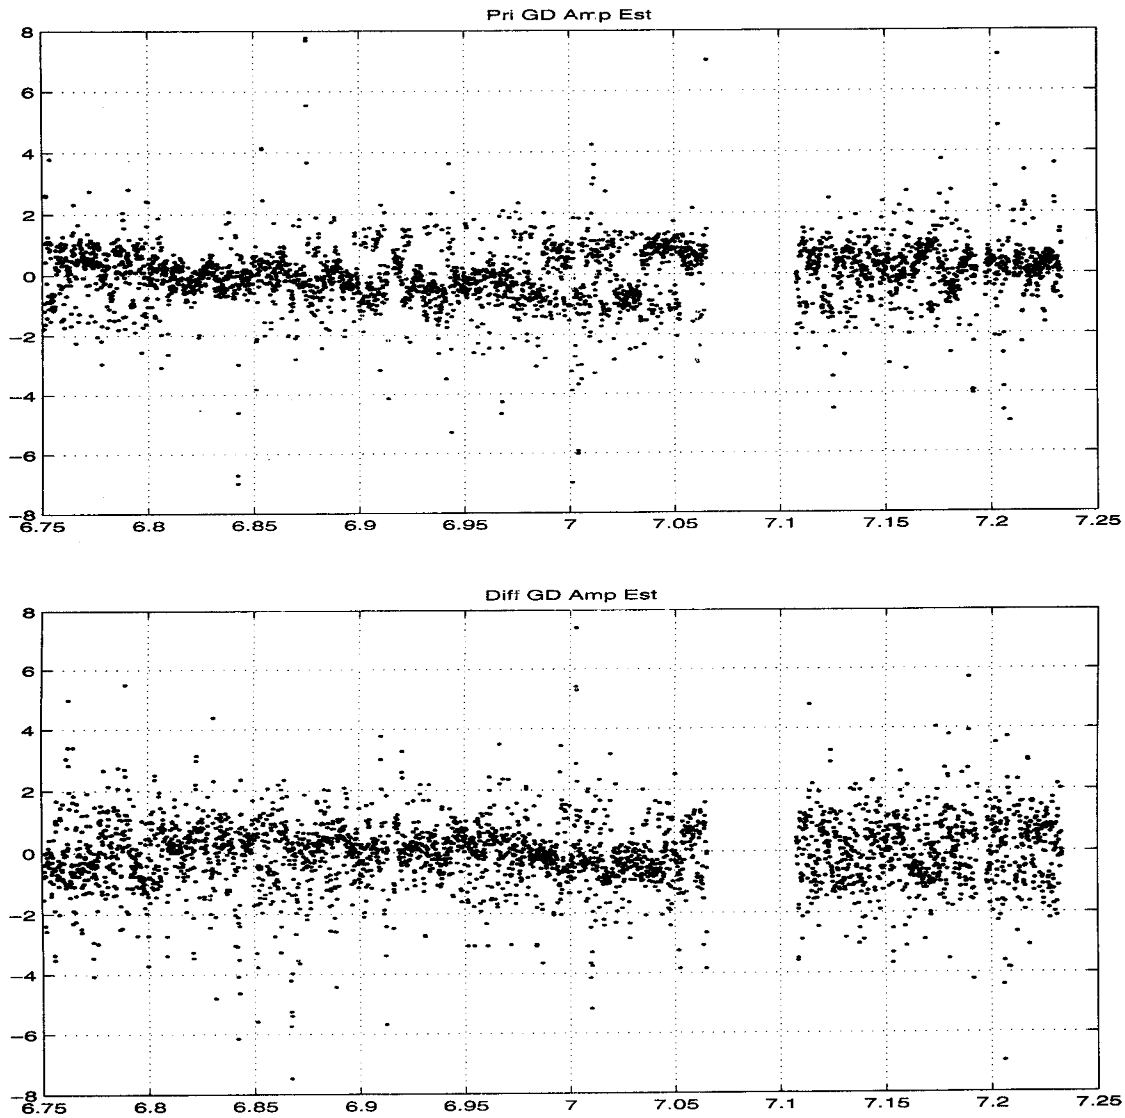

Figure 2. Amplitude group-delay estimates for the primary (top) and secondary (middle). This data was used to estimate the time dependent phase $\phi(t)$ prior to reprocessing the data.

$$\sigma_{gd} = \frac{1/\Delta\kappa}{\mathrm{SNR}(s)}. \tag{5}$$

It follows that delay estimation by phase tracking provides an rms error smaller than group-delay tracking by a factor of

$$\frac{\sigma_{\phi d}}{\sigma_{gd}} = \frac{\lambda_{wl}}{1/\Delta\kappa}, \tag{6}$$

(assuming that we can ignore the hopping back and forth from one fringe to another). For PTI, with an observation bandwidth of 2.0–2.4 μm and a mean observing wavelength of $\lambda_{wl} = 2.2$ μm, we have therefore $1/\Delta\kappa = 12$ μm and a variance of phase estimates that should be 5.5 times smaller than that of group delay estimates. There would therefore be a significant gain in performance if a group-delay estimator could be formulated that would behave as a phase estimator.

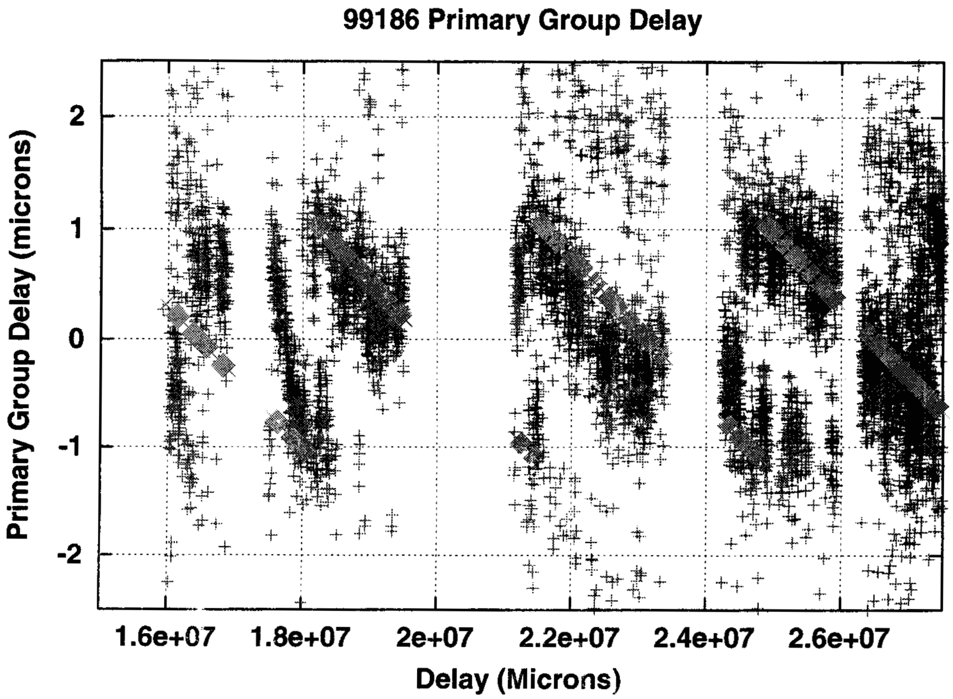

Figure 3. Comparison of the performance of the least-squares estimator with that of the normal group-delay estimator.

2. LEAST SQUARES ESTIMATION OF GROUP-DELAY USING PHASORS

When power-spectrum analysis is used for group-delay estimation, typically the delay corresponding to the peak is found by interpolating between the highest points in the power spectrum. The assumption implicit in this approach is that the phase of the Fourier transform at the peak is too noisy to be trusted, and so should simply be discarded. Whereas this would be true at low-light levels, at high light levels it provides a means for better estimates of group delay, with a precision approaching that of phase estimates.

If a zero-seeking phase tracker is used, we know that the phase at the central wavelength is always zero (modulo 2π), but that the slope of the phase as a function of wavelength could be almost any value. An estimate of the group delay will yield the number of fringes across the detector and a phase shift, which must be consistent with the phase being zero at the central wavelength.

The least-squares estimator can be described as follows. The data and its corresponding model, for sine and cosines is

$$f_k = C_0 \cos(2\pi\kappa_k x - \phi), \tag{7}$$

$$g_k = C_0 \sin(2\pi\kappa_k x - \phi), \tag{8}$$

where we have modeled a group delay x and an offset phase ϕ which may be due to the source or to the effects of dispersion (ie. if the white-light fringe is not actually centered on the peak of the fringe envelope).

The residuals to the fit can be written as

$$\sum_{k=1}^{K} \left[[C_0 \cos(2\pi\kappa_k x + \phi) - f_k]^2 + [C_0 \sin(2\pi\kappa_k x + \phi) - g_k]^2 \right].$$

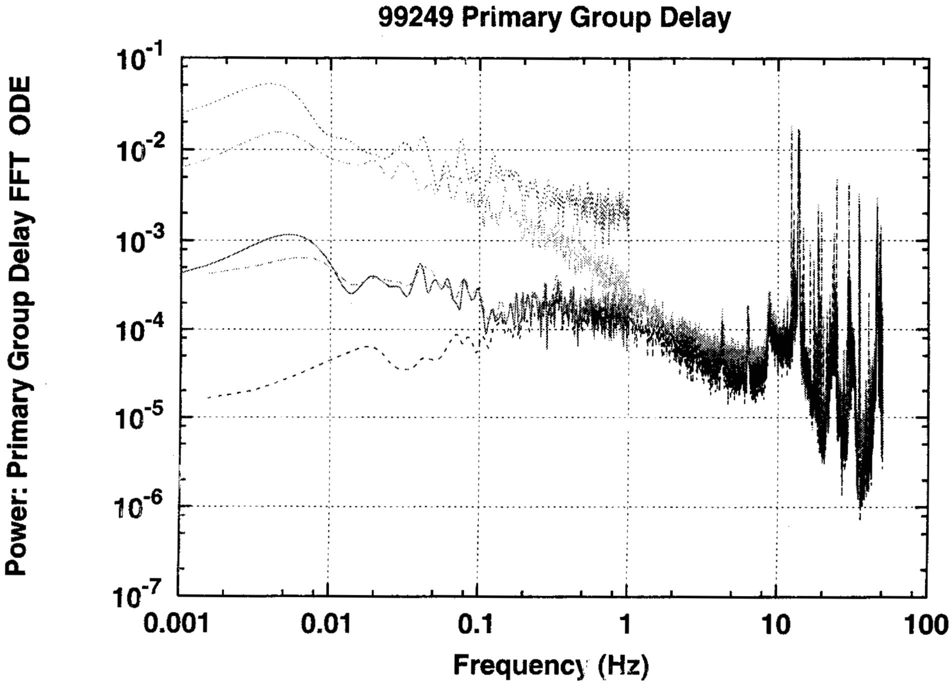

Figure 4. Power spectrum of the normal group-delay estimator and the least-squares estimator, shown along with the power spectrum of the 4-bin phase estimator used at PTI. One would expect that the least-squares estimator would show similar performance as the phase estimator and have an rms variation a factor of 6 smaller than the normal FFT estimate.

Expanding each of the terms and regrouping we have

$$\sum_{k=1}^{K} \left[C_0^2 + f_k^2 + g_k^2\right] - 2C_0 \sum_{k=1}^{K} \left[f_k \cos(2\pi\kappa_k x + \phi) + g_k \sin(2\pi\kappa_k x + \phi)\right].$$

If we wish to minimize the residuals we can ignore the series on the left side as it is constant; we must find the values of x and ϕ that maximize the series on the right. Therefore, if we wish to minimize the residuals we must find the values of x and ϕ that maximize the following series:

$$2C_0 \sum_{k=1}^{K} \left[f_k \cos(2\pi\kappa_k x - \phi) + g_k \sin(2\pi\kappa_k x - \phi)\right].$$

This least-squares estimator is sensitive to fringe phase, whereas our usual amplitude group-delay estimator is sensitive only to phase-slope (derivative of phase as a function of wavenumber).

Two passes through the data are needed for the data reduction:

1. First to determine the dispersion introduced by the air delay line, the smoothly-varying systematic difference between the group-delay and phase-delay — as a function of the total delay term, which is the same for primary and secondary, plus the systematic offset for primary and secondary.

2. Second to re-run the least-squares group-delay estimator using the proper model.

The new model would look like $\phi = \phi_{\text{offset}} + \text{const} \times \text{total delay}$.

3. RESULTS: OBSERVATIONS AT PTI ON 98193

The least-squares algorithm was included in vis, the standard data reduction program at PTI and primary-primary test data was processed, most notably the night of 98193. Although an improvement was expected in accuracy of a factor of 10 for data with a signal-to-noise ratio greater than 3, this was not seen. The histogram of the estimator showed several sets of peaks at equal spacings in delay, with the central peak sometimes offset from the peak located by the normal group-delay estimator. The width of the peak of the optimal estimator, was perhaps only a third as narrow as the GD estimator.

The following items are understood to be true:

- The usual, or *amplitude*, group delay estimate tracks the peak of the fringe envelope.
- The real-time fringe tracker tracks a position of constant phase.
- The peak of the envelope and the constant-phase location will slowly drift apart with time, and will wrap around after 4 meters of delay-line motion .
- The least-squares estimator is a phase estimator.
- The CT metrology measures a phase that is correctly unwrapped.
- The CT metrology follows fringe hops introduced by the fringe tracker and the changes in group delay that are made 0.5 seconds (based on the amplitude group-delay estimator calculations).
- The fringe trackers on each table are independent.

4. DISCUSSION

4.1. Does atmospheric turbulence degrade the quality of the least-squares fit?

When a fringe tracker is in lock, the delay line position *as a function of time* is determined by the siderial motion of the star, path-fluctuations due to atmospheric turbulence, and random fringe hops of one or more white-light wavelengths due to visibility fluctuations.

The group-delay *as a function of delay-line position*, depends on the different dispersive pathlengths that are encountered in each arm of the interferometer. Atmospheric turbulence adds changing dispersive pathlengths that are independent of the delay-line position. These path-variations are on the order of 1 μm rms for every meter of baseline, assuming an infinite outer scale (Tango and Twiss 1980, Eq. 4.2). So we have possibly 0.1 mm rms of atmospheric delay introduced over the baseline PTI typically uses.

Atmospheric turbulence adds a negligible amount of noise to our least-squares fit. If the fringe servo is locked, the group-delay will change by one wavelength for about every 4 meters of delay — and then it will wrap. That implies that the 0.1 mm rms of delay introduced by the atmosphere, introduces noise in the group-delay at a level of about 10^{-4} wavelengths rms. The least-squares fit is typically done using data for which the delay has changed by several meters, and so the $\sim$100 μm of atmospheric delay is entirely negligible.

4.2. Why do the amplitude group-delay estimates look so noisy?

The amplitude group-delay estimator usually works with noisy data, has relatively poor resolution, and often there are instances where the estimator fails altogether, producing outliers scattered at random locations in delay. When the amplitude estimator is performing well, there may yet be fringe hops due to visibility fluctuations. Multiple additive fringe hops also serve to spread the estimated group-delay across many wavelengths. If fringe hops are present and the estimator is otherwise performing well, the group-delay estimates may appear as sets of two or more parallel lines.

If we were observing a very bright star and the atmosphere were perfectly stable, the group-delay as a function of delay-line position would be a straight line with a slope determined by the dispersion. There would be very few outliers and no fringe hops — except for a single hop every four meters of delay when the group delay wraps to the other side of the fringe envelope.

5. CONCLUSION

The least-squares estimator of group-delay should provide delay estimation with the same accuracy as a phase estimator. It has been shown to work well under favourable conditions – high light levels, high fringe visibility, and good seeing conditions – as would be expected. There is as yet an unresolved difference between the its potential accuracy and the measured rms variations seen with laboratory experiments at PTI. There has also been considerable difficult in reliably predicting the dispersion one would expect as a function of delay, and it seems possible that the effects of atmospheric water vapor have not yet been fully understood.

For this implementation to be successful, we must be able to predict the longitudinal dispersion introduced by the air delay lines as well as any differential contribution due to atmospheric water vapour. The effects of water vapor are of particular concern, as it appears they may introduce a random fluctuation of more than one wavelength which obscures the linear term of the delay line. If this is the case then it may prove impossible to predict and subtract the dispersion terms. There are strong indications that the group-delay refractive index of air is very poorly modelled at near-infrared wavelengths. During the 1999 season, only half a dozen nights produced trends in group-delay that were readilly apparent in the data.

ACKNOWLEDGMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

J.T. Armstrong, D. Mozurkewich, L.J. Rickard, D.J. Hutter, J.A. Benson, P.F. Bowers, N.M. Elias, C.A. Hummel, K.J. Johnston, D.F. Buscher, J.H. Clark, L. Ha, L.-C. Ling, N.M. White, and R.S. Simon, "The Navy Prototype Optical Interferometer," Astrophys. J. **496**, 550–571 (1998).

M.M. Colavita, "Fringe visibility estimators for the Palomar Testbed Interferometer," Pub. Astron. Soc. Pac. **111**, 111–117 (1999).

M.M. Colavita, J.K. Wallance, B.E. Hines, Y. Gursel, F. Malbet, D.L. Palmer, X.P. Pan, M. Shao, J.W. Yu, A.F. Boden, P.J. Dumont, J. Gubler, C.D. Koresko, S.R. Kulkarni, B.F. Lane, D.W. Mobley, and G.T. van Belle, "The Palomar Testbed Interferometer," Astrophys. J. **510**, 505–521 (1999).

G. Daigne and J.-F. Lestrade, "Astrometric optical interferometry with non-evacuated delay lines," Astron. Astrophys. Supp. Ser. **138**, 355–368 (1999).

P.R. Lawson, "Group-delay tracking in optical stellar interferometry with the fast Fourier transform," J. Opt. Soc. Am. A **12**, 366–374 (1995).

W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C* (Cambridge, UK: Cambridge University Press, 1992).

J. Wyant, "Use of an ac heterodyne lateral shear interferometer with real-time wavefront correction systems," Appl. Opt. **14**, 2622 (1975).